

Attorney Docket No.: 021756-002100US

**PATENT APPLICATION**  
**METHOD OF AND SYSTEM FOR CREATING QUERIES THAT**  
**OPERATE ON UNSTRUCTURED DATA STORED IN A DATABASE**

Inventor(s): Srikanth Karimisetty,  
A Citizen of India,  
112 Fields Lane  
Peekskill, NY 10566  
  
Srinivasulu Puri,  
A Citizen of India,  
16 Lake Street, Apt. #7B  
White Plains, NY 10603  
  
Charanjeet Singh,  
A Citizen of India,  
809 Salt Court  
Redwood City, CA 94065  
  
Ravindra Akella,  
A Citizen of India,  
39 South Road  
White Plains, NY 10603  
  
Savita Durgada,  
A Citizen of India,  
163 Carrollwood Drive  
Tarrytown, NY 10591

Assignee: ORACLE INTERNATIONAL CORPORATION  
500 Oracle Parkway,  
Redwood City, CA, 94065  
A California corporation

Entity: Large

TOWNSEND and TOWNSEND and CREW LLP  
Two Embarcadero Center, 8<sup>th</sup> Floor  
San Francisco, California 94111-3834  
Tel: 650-326-2400

Oracle Ref. No.: OID-2003-067-01

## **METHOD OF AND SYSTEM FOR CREATING QUERIES THAT OPERATE ON UNSTRUCTURED DATA STORED IN A DATABASE**

### **CROSS-REFERENCES TO RELATED APPLICATIONS**

[0001] This application claims the benefit of U.S. Provisional Application No.

5 \_\_\_\_\_, entitled eSIGNATURES AND eRECORDS SYSTEM, by Srikanth  
Karimisetty et al. (Attorney Docket No. 021756-001901US), filed November 18, 2003,  
the disclosure of which is incorporated herein by reference. This application is also  
being filed concurrently with U.S. Application No. \_\_\_\_\_, entitled METHOD  
AND SYSTEM FOR ASSOCIATING AN ELECTRONIC SIGNATURE WITH AN  
10 ELECTRONIC RECORD, by Srikanth Karimisetty et al. (Attorney Docket No.  
021756-001900US); and with U.S. Application No. \_\_\_\_\_, entitled METHOD  
OF AN SYSTEM FOR SEARCHING UNSTRUCTURED DATA STORED IN A  
DATABASE, by Srikanth Karimisetty et al. (Attorney Docket No. 021756-002000US);  
and with U.S. Application No. \_\_\_\_\_, entitled METHOD OF AND SYSTEM FOR  
15 COMMITTING A TRANSACTION TO A DATABASE, by Srikanth Karimisetty et al.  
(Attorney Docket No. 021756-005100US); and with U.S. Application No.  
\_\_\_\_\_, entitled METHOD OF AND SYSTEM FOR COLLECTING AN  
ELECTRONIC SIGNATURE FOR AN ELECTRONIC RECORD STORED IN A  
DATABASE, by Srikanth Karimisetty et al. (Attorney Docket No. 021756-005200US);  
20 and with U.S. Application No. \_\_\_\_\_, entitled METHOD OF AND SYSTEM  
FOR DETERMINING IF AN ELECTRONIC SIGNATURE IS NECESSARY IN  
ORDER TO COMMIT A TRANSACTION TO A DATABASE, by Srinivasulu Puri et  
al. (Attorney Docket No. 021756-005300US), the disclosures of which are incorporated  
herein by reference.

25

### **BACKGROUND OF THE INVENTION**

[0002] The present invention relates to a method of and system for securely storing  
and accessing electronic records, capturing electronic signatures and securely  
associating captured electronic signatures with corresponding electronic records.

30 Embodiments of the invention are useful to a variety of companies in a variety of  
different industries and some embodiments are particularly useful in helping  
pharmaceutical, medical device, and food manufacturing companies ensure compliance

with Good Manufacturing Practice regulations (GMPs) as the companies produce and test products that people/animals use.

[0003] Historically, organizations manually tracked huge volumes of data from day-to-day business transactions in hard copy format. The requirements of maintaining such hard copies differed from one organization to another and in some instances depended in part on government led mandates for standard operating procedures. For example, if a pharmaceutical company based in the United States wanted to maintain records documenting the creation of a new drug from its inception to the clinical test stage it had to maintain a huge volume of paper based records in accordance with Food and Drug Administration (FDA) regulations. It can be appreciated that searches through such records for particular pieces of information could be a time consuming activity. It can also be appreciated that if access to particular classes of the hard copy records needs to be controlled for security reasons, physically separating the confidential documents from the regular ones becomes a highly tedious job.

[0004] Accordingly, many organizations have switched from paper-based records systems to electronic-based records (sometime referred to herein as eRecords). With the advent of eRecords, all business transaction records of an organization are stored electronically in a common data store of a business application software system. Some systems that store electronic records use XML (eXtensible Markup Language) technology because XML is an open, extensible and nonproprietary format.

[0005] While storing XML-based eRecords may significantly reduce paperwork previously required, the ability to efficiently search all existing eRecords is still a challenging task due to the flexible structure of the XML records themselves. The flexible XML structure of eRecords also makes restricting access to certain XML records in an organization a challenging proposition.

[0006] In addition to these challenges, eRecord systems must also be able to support the ability to electronically sign the electronic documents, ensuring that the appropriate personnel have reviewed and approved them. GMPs generally require signatures on transactions that affect product quality. Companies may also require signatures when moving the custody of goods from one location or department to another or when moving responsibility for manufacturing from one department to another. In general, wherever companies have generally required a paper signature in the past for such transactions, a signature is needed on the electronic document that replaces it once the company has made the switch from a paper-based world to electronic documents.

[0007] Accordingly, systems and methods for addressing the challenges in an electronic-based records system and systems and methods for handling electronic signature requirements in an electronic-based records system are needed.

5

## BRIEF SUMMARY OF THE INVENTION

[0008] Embodiments of the invention provide an improved electronic-based records system that addresses the challenges discussed above.

[0009] According to one embodiment of the invention, a method of collecting an electronic signature for an electronic record stored in a database is disclosed. The method comprises automatically creating an electronic record from data stored in a plurality of different database tables in response to the occurrence of a predetermined event; storing an instance of the electronic record in a common repository of electronic records that provides an audit trail that cannot be altered or disabled by users of the system; executing a rule associated with the electronic record to determine whether an electronic signature is required to connote review and/or approval of the electronic record; and if execution of the rule results in a determination that an electronic signature is required, marking the instance of the electronic record as unsigned and initiating a request to collect the required electronic signature. In one particular implementation of this embodiment, the electronic record is stored in a common repository of electronic records that provides an audit trail that cannot be altered or disabled by users of the database. In some implementations the electronic record is stored as unstructured data in a character large object (CLOB) format and the unstructured data comprises a well-formed XML document stored within a column of a table stored in the database.

[0010] In another embodiment, a method of associating an electronic signature with an electronic record is disclosed. The method comprises allowing a user to define an event that, upon occurrence, generates an electronic record that requires an electronic signature; allowing a user to define the fields stored in the electronic record; allowing a user to generate a map that maps data from underlying database tables to at least some of the fields defined for the electronic record; allowing a user to define a layout for displaying data in the electronic record on a computer display when an electronic signature for the data record is collected; allowing a user to identify a signatory approver for the electronic record; in response to the occurrence of the event, generating the electronic record and displaying the electronic record to the signatory

approver according to the defined layout; receiving an electronic signature from the signatory approver; and associating the electronic signature with the electronic record. In some implementations of this embodiment the method further comprises verifying the electronic signature prior to associating the signature with the electronic record.

5 [0011] In another embodiment, a method of committing a transaction to a database is provided. The method comprises initiating a database transaction; creating an electronic record that includes transaction data from the database transaction; executing a rule associated with the record to determine whether an electronic signature is required to connote review and/or approval of the electronic record, and if execution of  
10 the rule results in a determination that an electronic signature is required, requesting the electronic signature prior to committing the transaction to the database.

[0012] Another embodiment of the invention pertains to a method of searching unstructured data stored in a database. The method comprises storing a plurality of electronic records, each comprising unstructured data stored in a character large-object  
15 (CLOB) format in a column of a table of the database, in a common repository of electronic records in the database that provides an audit trail that cannot be altered or disabled by users of the system; creating a security protocol that protects the electronic records against unauthorized access; creating a query designed to identify electronic records in the database that meet criteria designated in the query; prior to executing the  
20 query, modifying the query in accordance with the security protocol to create a modified query; running the modified query against the unstructured data. In some implementations of these embodiments, access to records in the database is automatically granted unless security protocol restricts such access and the security protocol comprises a plurality of security rules that restricts access to the records within  
25 the database. In other implementations, access to records in the database is automatically denied unless security protocol grants such access and the security protocol comprises a plurality of security rules that grant access to the record within the database. In one particular implementation of this embodiment, the database is a common repository of electronic records that are generated from multiple data sources.  
30 In some implementations the unstructured data comprises a well-formed XML document stored within a column of a table stored in the database.

[0013] Another embodiment of the invention pertains to a method of searching unstructured data stored in a database where the method comprising storing unstructured data in a column of a database table; allowing a user to identify elements

in the unstructured data as indexed elements; creating an intermediate index into the unstructured data from the identified elements; and allowing a user to create queries on the unstructured data using the indexed elements. In some implementations of this embodiment the unstructured data comprises a well-formed XML document stored within a column of a database table. Also, in some implementations the unstructured data is part of an electronic record stored in a common repository of electronic records that provides an audit trail that cannot be altered or disabled by users of the database.

[0014] In still another embodiment of the invention, a method of intercepting a transaction instantiated by a database application to determine if an electronic signature is necessary to commit the transaction to the database is disclosed. The method comprises calling an application program interface to raise an event in response to a triggering action generated by the database application; initiating a workflow process that executes a rule to determine if an electronic signature is required to approve the transaction; and if execution of the rule results in a determination that an electronic signature is required for the transaction, instantiating a signature collection process. In some implementations of this embodiment the method further comprises obtaining an electronic signature in response to the signature collection process and thereafter, verifying the electronic signature and updating a file of the electronic record to indicate a valid signature was collected if the electronic signature is verified.

[0015] Other embodiments of the invention include computer systems comprising a processor, a database and a computer-readable memory coupled to the processor, where the computer-readable memory is configured to store a computer program that allows the processor to perform the methods described herein. And still additional embodiments of the invention are directed to computer programs stored on computer-readable storage mediums where the computer program comprises code for carrying out the methods described herein.

[0016] These and other embodiments of the invention along with many of its advantages and features are described in more detail in conjunction with the text below and the attached figures.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0017] Fig. 1A is a simplified block diagram of computer system 10 for managing electronic records and electronic signatures according to one embodiment of the present invention;

- [0018] Fig. 1B is a logical block diagram of system 10 shown in Fig. 1A according to one embodiment of the present invention;
- [0019] Fig. 2 illustrates one example of the data format used to store electronic records in evidence store 30 shown in Fig. 1B;
- 5 [0020] Fig. 3 is a more detailed block diagram of the system shown in Fig. 1B;
- [0021] Fig. 4 is a flow chart that depicts various steps involved with setting up system 10 to associate electronic signatures with particular electronic records according to one embodiment of the invention;
- [0022] Fig. 5 illustrates one example of an input form that can be completed by a user  
10 to enter information related to an event;
- [0023] Fig. 6 illustrates one example of an approval matrix that can be defined during the process set forth in Fig. 4;
- [0024] Figs. 7, 8 and 13-15 are examples of computer display screen shots generated by a graphical user interface to assist a user in creating indexed elements that can be  
15 used to create security rules and queries on electronic records stored in evidence store 30 according to one embodiment of the invention;
- [0025] Fig. 9 illustrates one example of database tables that can be used to track indexed elements as part of index 33 shown in Fig. 1B according to one embodiment of the invention;
- 20 [0026] Figs. 10, 12 and 20 are examples of computer display screen shots generated by a graphical user interface to assist a user in creating security rules according to one embodiment of the invention;
- [0027] Fig. 11 illustrates one example of a database table used to track security rules according to one embodiment of the invention; and
- 25 [0028] Figs. 13, 14, 17-19 and 21 are examples of computer display screen shots generated by a graphical user interface to assist a user in creating queries according to one embodiment of the invention.

#### DETAILED DESCRIPTION OF THE INVENTION

- 30 [0029] In some industries a process can be viewed as a pre-defined method of producing goods. Such processes typically include several events and sub-events where an event is an operation or group of operations to be performed to accomplish a task. The execution of events such as these are time, material and resource sensitive. Organizations often desire to track the execution of events to make sure that the event

has been completed as required by the process that the event is part of. To achieve this, checkpoints can be implemented at each event or sub-event to keep track of information such as: what is the event? who initiated the event? when was it initiated? who authorized the event? when was it completed? who confirmed the event completion? etc.

[0030] This information can be captured electronically and stored in a database so that it can be subsequently retrieved using query based user interfaces or reports. Embodiments of the invention allow a company or other organization to compile and store electronic records that track various events defined by the company.

10 Embodiments also allow electronic signatures to be captured and linked with their respective electronic data records so that the electronic signatures to be kept as part of a data record's audit trail. As used herein an "electronic signature" (sometimes referred to as an "eSignature") is a computer data compilation of any symbol or series of symbols executed, adopted or authorized by an individual to be the legally binding  
15 equivalent of the individual's handwritten signature. An electronic signature contains at least two distinct components: an ID and a password. An electronic signature connotes authorship, review and approval of data and can be displayed and printed with signed electronic records. A "digital signature" is an electronic signature based upon cryptographic methods of originator authentication, computer by using a set of rules  
20 and a set of parameters such that the identify of the signer and the integrity of the data can be verified.

[0031] Some embodiments of the invention operate in a closed system. As used herein, a "closed system" is an environment in which system access is controlled by persons who are responsible for the content of electronic records that are on the system.

25 In contrast, an open system is an environment in which system access is not controlled by persons who are responsible for the content of electronic records that are on the system. An example of an open system is an online trading community.

[0032] Fig. 1A is a simplified block diagram of computer system 10 for storing and updating electronic records, including electronic records associated with events,  
30 according to one embodiment of the present invention. As used herein an electronic record (sometimes referred to as an "eRecord") is a defined set of data captured from a moment in time by software. The data may include any combination of text, graphics, audio, pictorial or other information represented in digital form that is created, modified, maintained, archived, retrieved or distributed by a computer system.



Typically, some or all of the data is captured to the electronic record from multiple database tables. System 10 tracks electronic records in a manner such that recorded changes to the records do not obscure previously recorded information.

[0033] One example of electronic records includes records maintained by  
5 pharmaceutical, medical device, food manufacturing and other companies to ensure compliance with Good Manufacturing Practices (GMPs) and 21 CFR Part 11. GMPs are regulations that describe the methods, equipment, facilities, and controls required for producing human pharmaceutical products, veterinary products, biologically  
10 derived products, medical devices and processed food among other items. 21 CFR Part 11 establishes a uniform, baseline standard by which the Food and Drug Administration (FDA) will consider electronic records to be equivalent to paper records and electronic signatures to be equivalent to handwritten signatures. The baseline standard provides an enforceable mechanism for accepting electronic records and their associated  
15 signatures and provides a level of confidence that electronic records maintained in accordance with the rule will be of high integrity. While specific examples of the invention are often described below in conjunction with maintaining electronic records and signatures in accordance with GMPs and 21 CFR Part 11, it is to be understood that system 10 can be used to store and update electronic records in other industries that do not require conformance to GMPs and/or conformance to 21 CFR Part 11.

[0034] As shown in Fig. 1A, computer system 10 includes one or more processors 1  
20 that communicate with a number of peripheral devices via a bus subsystem 2. These peripheral devices may include a storage subsystem 3, comprising a memory subsystem 4 and a file storage subsystem 5, user interface input devices 6, user interface output devices 7, and a network interface subsystem 8. The input and output devices allow  
25 user interaction with computer system 10. A user may be a human user, a device, a process, another computer, and the like. Network interface subsystem 8 provides an interface to other computer systems and communication networks including communication network 9, which may be, for example, a local area network (LAN), wide area network (WAN) and/or public network such as the Internet. While not  
30 shown in Fig. 1A, in some embodiments network interface subsystem 8 is connected to one or more network server computers, for example by a LAN within network 9. The server computers connect computer system 10 to the Internet and provide appropriate security functions such as a firewall and the like.

**[0035]** Bus subsystem 2 provides a mechanism for letting the various components and subsystems of computer system 10 communicate with each other as intended. The various subsystems and components of computer system 10 need not be at the same physical location but may be distributed at various locations within network 9.

5 Although bus subsystem 2 is shown schematically as a single bus, alternative embodiments of the bus subsystem may utilize multiple buses.

**[0036]** User interface input devices 6 may include a keyboard, printing devices, a mouse, trackball, touchpad, a graphics tablet, a scanner, a barcode scanner, a touch screen incorporated into the display, audio input devices such as voice recognition systems, microphones, and other types of input devices. User interface output devices 10 7 may include a display subsystem, a printer, a fax machine, or non-visual displays such as audio output devices. The display subsystem may be a cathode ray tube (CRT), a flat-panel device such as a liquid crystal display (LCD), or a projection device. In general, use of the term “output device” is intended to include all possible types of 15 devices and ways to output information from computer system 10.

**[0037]** Storage subsystem 3 may be configured to store the basic programming and data constructs that provide the functionality of the computer system and of the present invention. For example, according to an embodiment of the present invention, software modules implementing the functionality of the present invention may be stored in 20 storage subsystem 3. These software modules may be executed by processor(s) 1. In a distributed environment, portions of the software modules may be stored on a plurality of computer systems and executed by processors of the plurality of computer systems. Storage subsystem 3 may also provide a repository for storing the various databases and files used by the present invention including evidence store 30 and database 34 25 shown in and discussed with respect to Fig. 1B.

**[0038]** Memory subsystem 4 may include a number of memories including a main random access memory (RAM) 4a for storage of instructions and data during program execution (e.g., execution of an instance of a software system for a database) and a read only memory (ROM) 4b in which fixed instructions are stored. File storage subsystem 30 5 provides persistent (non-volatile) storage for program and data files, and may include one or more hard disk drives, appropriate removable media cartridges, and other like storage media. One or more of the storage devices may be located at remote locations on other connected computers.

[0039] Computer system 10 itself can be of varying types including a personal computer, a portable computer, a workstation or any other data processing system. Due to the ever-changing nature of computers and networks, the description of computer system 10 depicted in Fig. 1A is intended only as a specific example for purposes of illustrating one embodiment of the computer system. Many other configurations of a computer system are possible having more or fewer components than the computer system depicted in Fig. 1A. For example, several other subsystems may be included in computer system 10 depending upon the functions performed by system 10.

[0040] Fig. 1B is a logical block diagram of computer system 10 according to one embodiment of the invention. As shown in Fig. 1B, system 10 includes a number of different logical components that are hosted on a computer system, which may include a number of distributed server systems and client systems. The logical components, which for illustrative purposes are shown as electronic record and electronic signature (ERES) setup component 20, ERES framework engine 22, security rule engine 24 and query engine 26 are executed by users 12 to allow system 10 to create, modify, maintain and search electronic records.

[0041] System 10 also includes an evidence store 30, a user interface 32 that allows users to interact with the logical components and the evidence store and a database 34. In some embodiments user interface 32 also allows users to interact with one or more business applications 36 that have the capability to invoke ERES engine 32 to generate and send electronic records to evidence store 30 as described in detail below. Business applications 36 can be any application used by an organization that accesses electronic records and/or allows users to electronically sign electronic records. Suitable examples of a business application 36 include any of the more than 160 software modules for financial management, supply chain management, manufacturing, project systems, human resources and customer relationship management that integrate into the Oracle Applications Suite developed and marketed by Oracle, assignee of the present invention.

[0042] Database 34 is a collection of information organized in a way that system 10 can quickly select desired pieces of data. Database 34 includes a database management system that enables data to be stored, modified and extracted from the database. The database management system within database 34 provides a database application program interface (API) through which database applications, such as business application 36, can access data managed by the database. Data managed by database

34 can be stored in relational tables that can be queried using any database language supported by the database including, for example, the popular structured query language (SQL).

5 [0043] In one embodiment, database 34 also allows applications to access data stored in the database by translating operating system I/O commands into database commands and/or through an alternative API referred to herein as the database file API. The database file API supports file operations similar to those supported by conventional OS file APIs but, unlike OS file APIs, incorporates the database API concept of transactions. That is, the database file API allows applications to specify that a set of  
10 file operations are to be performed as an atomic unit.

[0044] The database file API commands are translated to database commands by a database file server in the database management system. According to one embodiment database file server is object oriented. Thus, routines supplied by the database file server are invoked by instantiating an object and calling methods  
15 associated with the object. In one implementation the database file server defines a "transaction" object class that includes the following methods: insert, save, delete, update, commit and roll-back. The database file API provides an interface that allows external entities, such as application 36, to instantiate and use the transaction object class.

20 [0045] The methods invoked on a single transaction object may involve multiple file operations. If one of the operations fails during the instantiation of the transaction object, a roll-back process is invoked that undoes the changes made by the transaction object. In one particular embodiment, database 34 is a collection of software programs and/or components available from Oracle, such as Oracle Database. It is to be  
25 understood that database 34 can be any system, however, that organizes data in a manner that allows system 10 to quickly select desired pieces. It is also to be understood that components 20-36 are logically represented in Fig. 1B and that system 10 can include fewer, more or differently arranged logical components.

[0046] Evidence store 30 is a secure, common repository of electronic records that  
30 can be generated from multiple data sources, such as multiple database tables within database 34, in response to the occurrence of predefined events. As a single repository, all records that need to be kept for compliance with a companies business objectives (e.g., compliance with 21 CFR Part 11) are kept together. Thus, in order to search the

companies electronic records, there is no need to search multiple databases on multiple servers.

[0047] Generally, data is added to evidence store 30 but never deleted. Individual eRecords can be deleted in one sense (e.g., replaced by a newer record) but the original, deleted eRecord still remains part of the evidence store and can be retrieved if desired. In this manner, the evidence store is a secure database that maintains an audit trail for all records that are added, changed or deleted from the evidence store. In doing so, system 10 tracks for each eRecord when a change was made, who made the change and creates a record of the before and after information. The audit trail cannot be altered or disabled by users of the system and allows the evidence store to be reconstructed back to any given date and time.

[0048] Electronic records stored in evidence store 30 include unstructured data, such as data stored in character large object (CLOB) format, in one or more columns of a database table. In one embodiment the unstructured data includes a well-formed XML document stored within a single table or column of the database. While each XML document adheres to a structure (e.g., a particular DTD) and in one sense is thus structured data, the database column the XML document is stored in is unstructured in the sense that it can store XML data that adheres to a variety of DTDs and thus is not limited to storing data that adheres to a particular structure. The XML format provides portability and longevity to the data captured in electronic records. The XML document may include numerous proprietary and/or application-specific formats and is thus a highly flexible structure.

[0049] In one embodiment, evidence store 30 is part of database 34 and includes four different database tables that are separate from the database tables used to support the one or more business applications 36 and other functions of system 10. These four database tables include a table for storing eRecord details, a table for storing eRecord parameters, a table for storing eSignature details and a table for storing eSignature parameters. Separate tables are used to store eSignature and eRecord parameters because the parameters can be different for different eRecords and their number can also vary significantly. The separate tables maintain the parameters as name-value pairs so that additional columns in the main tables do not need to be added-an approach that is especially beneficial when it is not known how many parameters may exist.

[0050] Fig. 2 illustrates one example of database table formats in which eRecord details (table 30a) and eSignature details (table 30b) can be stored in evidence store 30

according to an embodiment of the invention. As shown in Fig. 2, table 30a includes a column for a primary key that is unique to each eRecord generated, a column that stores well-formed XML data generated for the record as unstructured data in CLOB format, a CLOB column that contains a copy of the XML data as it is displayed to users of

5 system 10 and various audit columns including columns for a timestamp when the eRecord was requested, a time stamp of when creation of the eRecord was initiated, the time zone the eRecord was created in, the status of the eRecord (e.g., pending, complete, error), the name of the entity or agent that requested creation of the eRecord, the event that the eRecord was created for, a field that tracks the number of times the

10 eRecord was printed or viewed and information tracking who created and last updated the eRecord. Table 30b stores a record of each eSignature requested and includes a primary key of a signature ID, a document ID that links the signature record to the eRecord (e.g., the primary key of the eRecord), the name of the user (signee) being asked to sign the eRecord, the signee's response, the date and time the eRecord was

15 signed, the time zone that it was signed in and the status of the eSignature (e.g., pending, complete, rejected, error). The signee's response may include, for example, details like the signing reason (whether it's a rework or first attempt), signer type (author, approver, reviewer, etc.).

[0051] Reference is now made to Fig. 3, which is more detailed logic-oriented block

20 diagram illustrating the interaction between selected logical components of system 10 in creating, updating and accessing data in the ERES evidence store according to one embodiment of the invention. The function-oriented diagram of Fig. 3 logically divides system 10 into four separate functions including, from left to right: eRecord and eSignature setup, eRecord and eSignature processing, eRecord and eSignature security

25 features, and eRecord and eSignature query features that operate to create, maintain and access evidence store 30. For convenience, each of these functional aspects of system 10 is discussed in more detail below.

### **Electronic Record and Electronic Signature Setup**

[0052] Setup component 20 of system 10 allows a user to perform all the various

30 functions that are required to integrate electronic record tracking and electronic signature capture alone or in conjunction with a business application 36. In the embodiment depicted in Fig. 3, setup component 20 includes multiple subcomponents

including a business event setup component 20a, an eRecord content setup component 20b and an eSignature rule setup component 20c.

[0053] Business event setup component 20a allows users to define events 40 that, upon occurrence of the event, trigger an action by another component of system 10, such as an action by ERES engine 22. As used herein, an event 40 is an occurrence in a computer application or program, such as an Internet or intranet application, that is significant to other objects in system 10 or to external agents. Event setup component 20a allows users to enable eRecord capture and/or eSignature collection upon the occurrence of the defined event by defining an ERES subscription 41 to the event. This can be done, for example, by storing two Boolean flags for each event: one flag denoting whether or not the event triggers generation of an eRecord and the other flag denoting whether or not the event triggers capture of an eSignature.

[0054] eRecord content setup component 20b allows users to define data type definitions (DTDs) 42, define mappings 43 between XML records and entries in database tables and create XSL (eXtensible Stylesheet Language) style sheets 44 that define layout settings for formatting and presenting eRecords to users. System 10 can store any defined database-to-XML mappings in a data repository, such as database 34 shown in Fig. 1B.

[0055] eSignature rule setup component 20c allows users to define rules 45 and an approval hierarchy 46 (approval matrix) for electronic signature events as well as store rule specific attributes such as the type of style sheet to be applied for a particular rule. Rules that can be defined by component 20c include rules related to how long system 10 will wait for obtaining signatures, whether multiple attempts to collect signatures will be made if necessary and, if so, the various timing issues associated with sending multiple attempts or reminder messages. Also in one embodiment of the invention, component 20c allows a user to select one of two distinct methods (rules) by which a required signature is collected: inline signature collection and offline (or deferred) signature collection. Details on these two collection methods are discussed later in the application and they are sometimes referred to as synchronous and asynchronous signature collection, respectively.

[0056] System 10 is typically set up by one or more users working with the various components of setup component 20 to define the various events system 10 tracks as well as the rules associated with tracking such events prior to using the system to capture, store and update electronic records. Fig. 4 is a flow chart that depicts various

steps involved with setting up system 10 to associate electronic signatures with particular electronic records according to one embodiment of the invention. The setup sequence includes three primary steps as shown in Fig. 4 including: defining events (step 50), defining event metadata (step 60) and defining a signature approval matrix (step 70). While these steps are depicted sequentially in Fig. 4, it is to be understood that the various setup steps may be performed partly or entirely in parallel and that they may occur in an order different than the order depicted in the Fig. 4.

[0057] As shown in Fig. 4, event definition step 50 includes a first substep 52 of defining the events 40 that require electronic signatures. In step 52 a user selects a unique name for each such event defined. The event can then be referenced across the framework of system 10 and outside of it by its chosen name. A user needs to know (or at least be able to identify) the event name in order to call the event with an appropriate application program interface (API) from one of business applications 36. In one embodiment the event name is a compound structure of identifiers separated by periods (.) as follows:

"oracle.apps.<product>.<component>.<object>.<event>"

where such a format allows users to organize the events defined into a classification hierarchy.

[0058] Fig. 5 illustrates one example of an input form 80 that can be completed by a user to enter information related to an event. As shown in Fig. 5, form 80 includes a field 81 to enter the event name, a field 82 that allows a user to select a name by which the event will be displayed in reports, a field 83 where a brief description of the event may be entered a field 84 that contains a status (e.g., enabled, disabled) of the event and a field 85 that indicates the application, such as a business application 36, to which the event belongs.

[0059] Once an event 40 is entered into the system by, for example, completing form 80 and selecting submit icon 86 shown in Fig. 5, a user can create a subscription 41 to the event by defining a triggering condition (Fig. 4, step 54) which evokes the event and specifying the processing that occurs when the event is evoked (i.e., when the triggering is met). The specified processing may include calling and executing a piece of code to perform a function and/or sending a message to a workflow process or other agent. A single event can have multiple subscriptions associated with it in which case the subscriptions are assigned a priority level that determines the order of execution. In



one particular embodiment the priority is set by a phase number stored in a column of a database table where lower phase number subscriptions are performed or initiated prior to higher phase number subscriptions.

[0060] One event subscription that can be defined by component 20a is the requirement for obtaining an electronic signature on an electronic record associated with the event (Fig. 4, step 56). A rule function created by a user with eSignature setup component 20c is associated with the subscription. The rule function determines if an eSignature is required and generates a snapshot of the data to be signed. The snapshot can be generated as an XML document in the rule function and displayed to required signees as dictated by XSL style sheet 44 defined in step 66.

[0061] System 10 can implement two types of eSignatures collection: synchronous (or inline collection) and asynchronous (or deferred collection). In one embodiment the phase of the eSignature subscription determines if the subscription is synchronous or not. For a synchronous signature subscription, the phase should be set to carry out eSignature collection as the initial action (e.g., by using the lowest possible phase number) to make sure that the eSignature subscription is the first subscription executed when an event occurs. In one embodiment, an asynchronous signature subscription can be created by setting the phase number above a phase number that represents a predetermined limit on the number of possible synchronous subscriptions. In other embodiments the synchronous/asynchronous nature of individual subscriptions can be set using other techniques, such as, for example, storing a Boolean attribute for each subscription.

[0062] Asynchronous signature collection allows system 10 to collect signatures when the signature does not need to be obtained immediately, the signer(s) are not in the same physical location or are otherwise not immediately available at the time of signature request or there are one or more items that the signer must verify prior to signing that would create a time lag between receipt of the signature request and the response. Asynchronous signature collection can be implemented with email or another messaging system. In contrast, inline signature collection allows system 10 to capture eSignatures when the signature needs to be captured immediately so as to not delay further processing. In one embodiment an inline signature collection process creates a pop-up window on the computer screen that displays the information captured in the eRecord for the signer to review. The signer must then review and sign the eRecord before further processing associated with the eRecord may occur.

[0063] Referring back to Fig. 4, the step of defining event metadata (step 60) includes defining a document type definition (DTD) 42 for each event 40 created in step 52 (step 62). DTD 42 determines the XML representation of all data that needs to be captured by the system for a particular event. The DTD can be defined in plain ASCII format using a text editor such as Notepad. The defined DTD 42 is then used in step 64 to map data, including data stored in database tables, to the event. Each event has certain database entities (objects) and attributes associated with it which need to be monitored and captured during event execution. Event data can span multiple entities.

[0064] Step 66 allows a user to decide what information in the DTD should be presented to a signee as well as how the information will be displayed when a signature is requested by an event by allowing the user to create an appropriate XSL style sheet 44. In step 68 DTD 42, XML mappings 43 and XSL style sheet 44 created in steps 62, 64 and 66, respectively, are stored in a computer-readable medium, such as a durable storage device used to store records for database 34 shown in Fig. 1.

[0065] eSignature rule setup component 44 allows users to define a signature approval matrix (step 70). The matrix includes rules that determine when an eSignature is required and/or when an eRecord is generated and if an approval is necessary, the identify of the necessary approver or approval group. Component 44 allows the matrix to define a single approver, multiple approvers where approval of each approver is necessary or a group of approvers where approval of only a subset of the group is necessary.

[0066] Fig. 6 illustrates one example of an approval matrix 90 that can be defined in step 72. Matrix 90 includes two rules 91 and 92 identified in the matrix as separate rows. Each rule has an associated attribute 93 and approver list 94. The attributes 93 used in matrix 90 include whether or not an eRecord is generated, whether or not an eSignature is required and if so, the style sheet used to present data for review to the approver. Rule 91 is executed if an event occurs at Plant 1 that uses the cyanide.

According to rule 91 all such events should be recorded as an eRecord and executed only upon obtaining approval of two individuals: James Dunn and Mark Arthur.

According to rule 92, all events occurring at Plant 1 should generate an eRecord for recording but unless other conditions are also met, e.g., the event includes cyanide, an eSignature is not necessary. If an eSignature is not required, users are not prompted to sign the transaction and in typical cases, the end user is shown a message that indicates the transaction is complete.

**Electronic Record and Electronic Signature Processing**

[0067] In operation, ERES framework engine 22 responds to events generated by business application 36 or other business transaction logic to create electronic records, collect electronic signatures and store electronic records and signatures in evidence store 30 as appropriate. ERES engine 22 can assign each instance of an event (each time the event is triggered) a unique event key that system 10 uses to track eRecord data associated with the particular event instance. The unique key can be passed as a parameter to an API that raised the event.

[0068] Referring again to Fig. 3, the blocks within ERES engine 22 represent method steps performed by engine 22 in response to the triggering of a business event 40 by a business application 36 that results in generation of an eRecord to be stored in the evidence store. In one embodiment, the event trigger process results in business application 36 calling an API that initiates a work flow process. The API may include, for example, the event name and an event ID, an indication of whether deferred signature collection mode is allowed for an event or not, the business application that the event was raised from, a username of a person that raised the event and any API or function that should be invoked by system 10 after the completion of the event among other information. The initiated work flow process is then run by the ERES engine.

[0069] As shown in Fig. 3, the ERES engine creates an eRecord in response to the trigger event (step 80). The eRecords are based upon the DTDs 42 and XML mappings 43 prepared under the guidance of eRecord content setup component 42 for the event. In one embodiment, upon creation, individual eRecords are stored in a database table that includes the XML data generated from mappings 43 as the mapped data existed at creation time, a copy of the XML data as it is displayed to users in the associated XSL style sheet, a primary key that is unique to each eRecord created in system 10, various time entries such as the timestamp when the eRecord was requested, the time stamp creation of the eRecord was initiated, the time zone the eRecord was created in, the status of the eRecord (e.g., pending, complete, error), the name of the entity or agent that requested creation of the eRecord, the event that the eRecord was created for and a field that tracks the number of times the eRecord was printed or viewed.

[0070] If a particular event requires an eSignature, ERES framework engine 22 evokes the appropriate signature collection rule 45 created under eSignature collection rule setup component 20c and instantiates a process to collect approval signatures according to the approval matrix 46 (step 81) associated with the rule. The instantiated

process routes notification that a signature is needed to each necessary approver (step 82). In one embodiment a record of each eSignature requested is stored in evidence store 30 in a database table that includes a primary key of a signature ID, a document ID that links the signature record to the eRecord (e.g., the primary key of the eRecord),  
5 the name of the user (signee) being asked to sign the eRecord, the signee's response, the date and time the eRecord was signed, the time zone that it was signed in and the status of the eSignature (e.g., pending, complete, rejected, error).

[0071] The notification of the eSignature request can be routed by, for example, email or a pop-up window. In either case, the requested signee is asked to review  
10 information associated with the eRecord 80 before granting approval or rejecting the request. The information presented to the signee is formatted according to the style sheet 44 that was created for such purpose (step 83). Once an eRecord has been completed and signed by each necessary signee, ERES engine 22 verifies the electronic signature and changes the status of the eSignature request to "complete" (step 84).  
15 Signature verification can be done, for example, by executing a function in the background to determine if all signatures have been obtained and comparing the captured username and password pair for each signature to the system's valid username/password pairs stored in database 34.

[0072] Events raised by business application 36 that create eRecords typically also  
20 result in an update to one or more of the tables of database 34 that tracks data associated with the business application. In this sense the business event can be thought of as initiating a transaction that may or may not get completed. Records of the transaction are stored in the evidence store regardless of whether or not the transaction is completed. i.e., the evidence store provides a complete history of a business event  
25 including failed or incomplete events. If the transaction initiated by the event is successfully completed, however, details of the transaction are committed to database 34. The actual transaction data that is updated in database 34 will vary according to the action initiated by the business application.

[0073] As one example, consider a scenario in which a trigger event is defined in one  
30 business application to generate an eRecord and an eSignature request in response to the creation of a purchase order (the purchase order is the transaction). Thus, when a user select a "create purchase order" option in this particular business application, an eRecord is automatically generated and a work flow process is initiated to capture the appropriate signature to authorize the purchase order. Upon its creation, the eRecord is

stored in evidence store 30 with all the various fields of the purchase order completed according to the data that was captured in the eRecord from the mappings 43. The eSignature collection process is then started. The signature collection process displays the purchase order with the captured data according to a predefined XSL style sheet to each signee. The eRecord stored in evidence store 30 stores the captured data in XML format and also stores the data as formatted in the eSignature request.

**[0074]** If authorization for the purchase order is approved by all necessary signees, each signee field in the eRecord in evidence store 30 is updated as "complete" and the purchase order is committed against database 34, i.e., the business transaction is completed. If authorization is denied by a signee with appropriate authority, the eRecord in evidence store 30 is updated to indicate which signee "rejected" the purchase order and the purchase order is not committed against database 34. This may require rolling back some portions of the purchase order transaction process that updated one or more portions of appropriate tables in database 34 so that database 34 is in the same state it was in prior to the purchase order request. In such a case, evidence store 30, however, will store all the necessary information that indicates a particular purchase order was created, authorization for the purchase order was obtained and ultimately denied by a particular individual.

**[0075]** In another example, the purchase order may include one or more fields that require data entered by a user through an electronic form-based data entry process. In such a form-based process, the eSignature request is not initiated until the user completes entry of data in the form and selects an icon such as "obtain authorization". At that time, the eRecord includes the data entered by the user along with any other data captured from database 34 that is part of the eRecord. For example, the automatically captured data may include information such as a customer name and a shipping address and the data entered via an electronic form may include a product code and a quantity indicated. The eRecord will store the purchase order in XML format including the captured customer name and shipping address and the entered product code and quantity.

**[0076]** After an event has been completed, system 10 next invokes the API that was identified in the original event API call, if an API was designated, for post-ERES processing. All eRecords generated and placed in evidence store 30 are committed to the evidence store 30 permanently, i.e., no eRecord can be deleted from the evidence store tables in the ordinary course of business. Also, the XML data captured in each

eRecord is captured as a permanent record that also cannot be changed in the ordinary course of business.

### **Electronic Record and Electronic Security Model**

[0077] As indicated above, the eRecords stored in evidence store 30 are a repository of very critical information for a company that often needs to be queried for various reasons ranging from internal users perusing the information to regulatory authorities inspecting process records. The information contained in these records can be confidential and critical to the nature of the business. Accordingly, embodiments of the invention allow a company to restrict access to eRecords to prevent any unauthorized access. Such security controls are defined and performed by security rule engine 24.

[0078] Security rule engine 24 allows users to identify and tag secure XML data records within evidence store 30, create security rules and, for users having the appropriate security privileges, turn selected security features of system 10 ON and OFF. Security rule engine 24 also provides security mechanisms to ensure that eRecords can never be removed from evidence store 30 and that electronic signatures cannot be excised, copied or otherwise transferred to falsify an electronic record by ordinary means. In some embodiments, security rule engine 24 also interfaces query engine 26 to evidence store 30 to ensure that queries on the evidence store do not take place if any active security rule indicates the user initiating the query does not have the necessary access rights.

[0079] Security rule engine 24 restricts access contingent on the content of an eRecord and the event for which the eRecord was created. For example, if a number of eRecords are created for a Lot Creation event, a responsible user (e.g., a Security Administrator) should be able to grant access to Lot Creation eRecords with particular lot numbers to individual users or groups of individual users. Engine 24 can be configured to operate in one of two modes: restrict mode in which access to eRecords is granted by default and users or responsibilities can be restricted as required and grant mode where access to eRecords is restricted by default and users or group basis based upon responsibilities assigned to the group are granted access to specific records based on the values or data contained in the record. Security rule engine 24 then allows security rules that grant or restrict access to secured content to be created on an individual basis or responsibilities.

[0080] The eRecords stored in evidence store 30 are XML documents adhering to particular DTDs defined by users of system 10. Before a security rule can be created using a particular XML element, the element has to be identified as a "secure element" (Fig. 3, function 90). A secure element is essentially an XML element identified in a particular DTD with a special use of being able to create security rules. Security rule engine 24 allows users to identify secure elements via a graphical user interface (GUI). In one embodiment, the first step in identifying a secure element is to select the element as an indexed element or an "IXE". Once an element is indexed, security rules and queries (discussed below) can be created for it.

10 [0081] System 10 can track indexed elements by setting an appropriate Boolean flag as an attribute of the element. That is, an element is either indexed or not indexed. In another embodiment the indexing of elements can be tracked by a character field that can have one of two values: e.g., I for indexed, N for not indexed. The default value for each element upon creation of a DTD is "not indexed". This default can be changed by a user using the GUI.

[0082] Fig. 7 illustrates one example of a screen shot (computer display page 100) from a GUI that can be employed to select elements to be indexed. As shown in page 100, XML elements can be searched on the basis of the element name (field 101) or their display name (field 102). A user also has the option of displaying all elements in system 10 that are capable of being indexed, all elements that are already indexed or all elements that are not indexed by selecting an appropriate choice in field 103. The results of the user's search are displayed in area 104.

[0083] Using icon 105 or field 106, a user can select to create an IXE or update the attributes of an element presented in results area 104 so that its attributes can be changed to an IXE or changed from an IXE to a non-indexed element. An example of an IXE creation screen shot (display page 110) is shown in Fig. 8. As shown in Fig. 8, page 110 includes selection boxes 111 and 112 that allow the user to designate uses for the IXE including whether the IXE is used for security purposes (a secure element), query purposes (a query element) or both (a secure and a query element). Page 110 also includes a field 113 for assigning the IXE a user-friendly name that will appear to users when the IXE is used to create a security rule or in a query.

[0084] System 10 creates a domain index (shown in Fig. 1B as index 33) from the indexed elements, or updates the existing domain index, by, for example, selecting apply icon 114 shown in Fig. 8. Index 33 facilitates identifying records containing the

elements in the evidence store and includes separate sections that are defined for each indexed element. In one embodiment, index 33 is created using Oracle's interMedia Text product and stored in the internal tables of database 34.

[0085] Fig. 9 illustrates one example of database tables that can be used to track indexed elements as part of index 33. Shown in Fig. 9 are three separate tables: table 33a which is used to store indexed elements for eRecords, table 33b which is the translation table for the indexed elements and table 33c which is used to store the usage of an indexed element. As shown in Fig. 9, each of tables 33a-33c include various audit columns that track when an element was created, who created it, etc. Table 33a also includes a indexed element primary key (`element_id`), a field that map the indexed element to an XML element in a particular DTD, fields that identify the entry in the index (section) for the indexed element as well as the unique tag associated with the entry and a field that tracks whether the indexed element is indeed indexed. Table 33b contains a field that matches the `element_id` to a display name that can be easily digested by a user and a field that tracks a description of the element. Finally, table 33c tracks the usage of the indexed element as a query element or secure element as discussed further below.

[0086] An element chosen to be an IXE is either a generic IXE or a DTD-specific IXE. A generic IXE is an IXE that is applicable to all eRecords in evidence store 30 irrespective of the DTDs. A DTD-specific IXE is an IXE that is applicable to eRecords based only on a particular DTD. This is the case when, for example, the same XML element has different semantics associated with it across different DTDs. Once an element is indexed and designated as a secure element, engine 22 allows users to create security rules with the element (Fig. 3, function 91).

[0087] A security rule creates a restriction or provides a grant similar to:

"Allow user James (user id: *JASDE*) to access eRecords for the event *Formula Approval* having value *Yeast* for *Formula Ingredient*"

where "Formula Ingredient" is the secure element for the event "Formula Approval" and access to value "Yeast" is being granted to user id "JASDE". Similarly, access to a particular user can also be restricted for a specific value. Engine 22 also allows security rules to be established for groups of users having common responsibilities or privileges. For example:



"Allow responsibility *Manager* to access all eRecords for the event *Employee Creation* having value *Salary* for *Employee Detail* "

[0088] Using such security rules, access to the eRecords for specific events can be restricted based on the contents of the eRecords; specifically the value of the secure elements in the eRecord. In addition to allowing users and responsibilities access or restricting such access based on an eRecord's content, engine 22 allows a group of users to be granted access to eRecords while restricting particular users within the group to the records based on their content and it allows a group of users to be restricted from accessing eRecords while allowing particular individuals within the group to be granted such access based on the record's content.

[0089] Security elements can be searched for in a manner similar to that described for indexed elements. Fig. 10 illustrates one example of a screen shot (computer display page 120) from a GUI that can be employed to find already defined security elements. As shown in page 120, security elements can be searched on the basis of the element name (field 121), event name (field 122) and either user (field 123) or responsibility (field 124). The results of the search (security rules created for the security element) are displayed in area 125.

[0090] As shown in Fig. 10, each security rule has a start date and an end date. This allows a user a lot of flexibility in enabling or disabling security rules without having to delete them or re-create them. A security rule can be deleted by selecting the trash icon 126 for the rule and a new security rule can be created by selecting create security rule icon 127.

[0091] Fig. 11 shows one example of a table 33d used to track security rules created according to some embodiments of the invention. Table 33d includes a primary key (rule\_id), a field that references the security element (element\_id) upon which the rule operates, various fields for tracking the parameters of the rule and audit columns that track who created the rule, when it was created and modified, etc.

[0092] An example of a security rule creation screen shot (display page 130) is shown in Fig. 12. As shown in Fig. 12, page 130 displays the name of the secure element for which the rule is being created (field 131), the event name (field 132) and allows a user to enter a value in field 133 that, if present in an eRecord, will result in the security condition (grant or restrict) as entered by the user in field 136 being applied to the record. The user creating the rule is preferably aware of whether security rule

engine 24 has been set to operate in high security mode (default of restricting access to all eRecords unless a rule explicitly grants a user/responsibility access to the record) or low security mode (default of granting access to all eRecords unless a rule explicitly restricts a user/responsibility from accessing record).

- 5 [0093] Table 1 below lists five exemplary rules and the effect such rules have if engine 24 is set to operate in high security default mode.

Examples of Rules vs. Effect in High Security Mode		
Rule No.	Security Rule	Effect
1	Grant access to user U1 for formula change eRecords having dry yeast.	None. Access already granted.
2	Grant access to responsibility R1 for formula change eRecords having dry yeast.	None. Access already granted.
3	Restrict access to responsibility R1 for formula change eRecords having dry yeast.	Access for R1 restricted.
4	Restrict access to user U1 for formula change eRecords having dry yeast.	Access for U1 restricted.
5	1 + 3	Access for U1 granted but rest of R1 restricted.

**TABLE 1**

- As evident from the Table 1, creating either rules 1 or 2 does not make much sense as the rules do not have any effect.
- 10

[0094] Table 2 below lists five exemplary rules and the effect such rules have if engine 24 is set to operate in low security mode.

Examples of Rules vs. Effect in Low Security Mode		
Rule No.	Security Rule	Effect
1	Grant access to user U1 for formula change eRecords having dry yeast.	Access granted to U1.
2	Grant access to responsibility R1 for formula change eRecords having dry yeast.	Access granted to U1.
3	Restrict access to responsibility R1 for formula change eRecords having dry yeast.	None. Access already restricted.

Examples of Rules vs. Effect in Low Security Mode		
Rule No.	Security Rule	Effect
4	Restrict access to user U1 for formula change eRecords having dry yeast.	None. Access already restricted.
5	2 + 4	Access for U1 granted but rest of R1 restricted.

TABLE 2

As evident from Table 2, creating rules 3 and 4 when engine 24 operates in low security mode does not make much sense because those rules have no added effect.

[0095] Referring to Fig. 12 again, the user or responsibility that the rule restricts or grants access to eRecord for can be identified in fields 134 and 135, respectively.

Finally, start and ending dates for the rule can be entered in fields 137 and 138. If no end date is entered, the rule will be in effect for perpetuity. If no specific start date is entered, the rule will take effect immediately (e.g., the current date will be entered by default).

[0096] A user with appropriate system level privileges can use security rule engine 24 to enable and disable the security features (Fig. 3, function 92). In one embodiment, when security is enabled queries are modified using the defined security rules before running the queries against the evidence store. Referring to Fig. 3, queries can be created by users 12 working with query engine 26 and user interface 32 in the manner explained below. Once a specific query 95 is created and selected to be executed against the evidence store, security engine 24 determines if there are any security rules that are relevant to the query. If there are, the query is dynamically modified by security rule engine 24 based on the defined security rules that are relevant to the query to create a modified query 96. The modified query is then executed against the evidence store and the results are returned to the user. If there are no relevant security rules, the original query (query 95) is run against the evidence store. In one embodiment the determination on whether or not to create a modified query 96 is made by identifying the indexed elements within original query 95 and identifying secure elements created from the indexed elements. A modified query 96 is generated if any secure elements have been created from the same indexed elements used in query 95.

[0097] In one embodiment, virtual private database (VPD) technology available from Oracle Corp., the assignee of the present application, is used to dynamically modify

each query using the defined security rules prior to running the query against the evidence store. Oracle's VPD technology allows security engine 24 to modify SQL query statements based on a WHERE condition (known as a predicate) returned by a function that implements the security policy. The SQL statement is modified

5 dynamically in a manner transparent to the user using any condition which can be expressed in, or returned by, a function.

[0098] As an example of one embodiment of the invention where evidence store 30  
evidence store 30 stores eRecords having XML data that include separate field  
INGREDIENT and DEV\_STAGE, consider a query generated by a user to identify all  
10 eRecords in pertain to ingredient "sugar" where a security rule exists that prevents  
eRecords associated with a "clinical stage" of the development process from being  
accessed. The original SQL format for such a query may be represented as:

```

15      select document_id
        from edr_psig_documents
       where contains (psig_xml, 'sugar within INGREDIENT')>0

```

where the eRecord is identified by its "document\_id", "edr\_psig\_documents" is the  
name of the table in evidence store 30 that stores eRecords and "psig\_xml" is the  
CLOB column of XML data within the edr\_psig\_documents table. Security engine 24  
modifies the original query by converting the security rule into a predicate that is  
20 AND'd with the original predicate of the SQL statement. The modified query may then  
be represented as:

```

25      select document_id
        from edr_psig_documents
       where contains (psig_xml, 'sugar within INGREDIENT')>0
         and contains (psig_xml, 'clinical_stage' within DEV_STAGE') +0=0

```

where the second contains clause was added by security engine 24.

[0099] As evident from the above described indexing process, a lot of flexibility is  
provided that allows different semantics to be associated with the same XML element  
or the same semantics to be associated with different elements. This in turn allows  
30 organizations to exercise a tremendous amount of flexibility in securing eRecords to  
obtain information from data. As an example of this flexibility, consider the following  
scenario. A user identifies that XML element ITEM\_NO in a particular eRecord for an  
Item Creation event refers to an identification number of a particular class of items and  
creates a secure element called Item Number using ITEM\_NO for these records. The  
35 user can also identify that XML element FORMULA\_INGREDIENT in an eRecord

that conforms to a different DTD for a Formula Creation event also actually refers to the same class of identification numbers as in the previous eRecord. The user can then create a secure element called Item Number using FORMULA\_INGREDIENT for these records and thereby link certain security rules defined by the user on the Item

5 Number secure element to both types of eRecords. The user can also identify that the XML element ITEM\_NO in a completely different set of eRecords adhering to a third DTD actually means something entirely different, such as a contract term in a Procurement Contract event. The user can then go ahead and create a secure element called Contract Term using the ITEM\_NO element for these eRecords. Based on this

10 scenario, a user can then create a security rule that, for example, restricts access to user "Neo" for all eRecords that include the secure element Item Number having a value of 057. This rule would then secure such eRecords generated from Item Creation and Formula Creation events but would not secure eRecords generated from a Procurement Contract event.

#### 15 **Electronic Record and Electronic Query Engine**

[0100] The eRecords stored in evidence store 30 often need to be queried for various reasons ranging from internal users perusing the information to regulatory authorities inspecting process records. Query engine 26 allows users to identify XML elements upon which queries and searches are to be performed and allows users to generate and

20 execute queries based on the identified elements. To perform such queries, however, a user needs to be able to distinguish fields in different DTDs defined by the users in step 62 (Fig. 4) that have identical XML field names yet represent different pieces of data. For example, evidence store 30 may include many different electronic records that have an XML document identifying an "item number" that is part of the record. The "item

25 number" in each XML document may adhere to different DTDs, however, and thus may refer to completely different item lists where a particular number equates to different items in the different DTDs.

[0101] As an example, in one DTD (DTD 1) that defines the ingredients for drug X, "item number" may be the unique identifier of each ingredient used to manufacture the

30 drug. In another DTD (DTD 2) that defines the results of a comprehensive drug study, "item number" may be a unique identifier that identifies which drug or placebo was given to a patient. Thus, item number 003 in DTD 1 may be "glycerin" while item number 003 in DTD 2 may be "placebo". Accordingly, a query in evidence store 30 for

all records that involve "glycerin" needs to be able to do more than search for an item number of 003.

[0102] Embodiments of the invention allow users to use query engine 26 to identify the XML elements of each DTD used to generate an eRecord that can be queried within  
 5 evidence store 30. As an initial step in defining query elements, the element has to be indexed as described above with respect to security rule engine 24. The indexing process builds a domain index (Fig. 1, index 33) on top of the tables that make up evidence store 30 as described above. The indexed elements can be designated as query elements (Fig. 3, function 93) using the process also described above and used in  
 10 queries. This process provides the same flexibility in using query elements that was described above with respect to secure elements.

[0103] In one embodiment, queries elements can be created under the guidance of a GUI using a display page 140 such as that shown in Fig. 13. As shown in Fig. 13, the query creation page 140 is divided into two sections: a top section 141 allows a user to  
 15 specify query criteria and a bottom section 142 displays the results of a given query. In top portion 141, a user can select between a simple query (tab 143 or an advanced query (tab 144). The query options in a simple query include event name (field 145), a date range (fields 146, 147) and a user id of an eRecord signer (field 148). Such queries can be run by selecting find icon 149.

20 [0104] Bottom portion 142 of screen 140 displays header level information from eRecords matching the query criteria including the business event name (field 150), the unique identifier of the event (field 151) and time-related information (fields 152, 153). If desired A user can select to print individual ones of the eRecords by clicking on print selection boxes 154.

25 [0105] An example of an advanced query screen page 160 is shown in Fig. 13. Page 160 allows a user to directly query the XML format of eRecords by allowing a user to select specific XML elements in fields 161 and enter particular values to be searched in field 163. Condition field 162 allows the user to select whether the query is for records that equal the value chosen in field 163, for records having a value less than the chosen  
 30 value, for records having a value greater than the chosen value, etc. Additionally, a user can include multiple selection criteria in a single query by adding another row to the query and selecting whether the row is to be logically AND'd or OR'd (field 164) with the previous criteria. The advanced query page returns query results to an area 165 that is similar to area 142 shown in Fig. 12 screen. Details of individual eRecords

can be displayed from either of query results sections 142 or 165 by clicking on an appropriate selection box 155 or 166, respectively.

[0106] Once a query is created (Fig. 3, query 95) it can then be executed against the evidence store (Fig. 3, function 94). Before doing so, however, the query is modified based on the security rules (Fig. 3, modified query 96) as explained in detail above. The modified query is then executed against evidence store 30 only if it does not violate any of the security rules.

#### **Example of Query Capabilities and Security Measures Provided by Invention**

[0107] Figs. 15-21 further illustrate the capabilities and functionality that can be obtained using some embodiments of the invention. Specifically, Figs. 15-21 represent exemplary screen shots and data formats that are created by system 10 in order to implement a scenario where a company that controls a manufacturing application that requires the creation of 'items' and 'lots' containing those items. The company desires that data for business transactions resulting in the creation of either of these two entities be captured as eRecords and stored in evidence store 30. Each eRecords includes a well-formed XML document stored in a single column (column name) of a table (table name) in evidence store 30.

[0108] In this exemplary scenario, a transaction that results in the creation of an item is referred to as an 'Item Creation' event and a transaction that results in the creation of a lot is referred to as a 'Lot Creation' event. Thus, eRecord are created for both Item Creation and Lot Creation events whenever such events take place. The XML documents (eRecords) that are created for each of these events adhere to different DTDs, however, as defined by a user using data setup component 20. In this scenario, the predefined DTDs for both Item Creation and Lot Creation events include an XML element labeled "ITEM\_NO". While this XML element functionally refers to an item number because of the way the DTDs are defined, in the case of an Item Creation event, end users identify the item number with an 'item' while in the case of a Lot Creation event end users identify the item number as a lot item as shown in Figs. 15 and 16, which are screen shots showing the creation of indexed elements that can be used as both query and secure elements from each XML element by a user. Fig. 17 is a screen shot that demonstrates the XML element 'item\_no' is associated with two different indexed elements after creation of the indexed elements as shown in Figs. 15 and 16.

[0109] Once the set up of the indexed, query and secure elements is complete, queries using query elements and security rules using secure elements can be created. Fig. 18 is a screen shot 170 illustrating one example of a query that can be created and executed on the data. As shown in Fig. 18, a query can be generated by entering a query element in field 171 and a value in field 172. Query elements can be selected in field 171 from a list of query elements presented in drop down menu format when find icon 173 is selected. In Fig. 18 a user has created a query that will return all eRecords that include a 'lot item' (previously defined as XML element 'item\_no' in a Lot Creation event per Fig. 14) that has a value of 4101. The results of the query are shown in area 174. Fig. 19 is another screen shot 180 illustrating another example of a query that can be created and executed on the data. As shown in Fig. 19, a user has created a query that will return all eRecords that include an "item" (previously defined as XML element 'item\_no' in an Item Creation event per Fig. 13) that has a value of 4101. The results of the query are shown in area 184. As evident from a comparison between Figs. 18 and 19, each query only returned results where the XML element 'item\_no' equaled 4101 for the DTD the query element is associated with. The XML element(s) that have the same value but that are associated with a different DTD were not returned.

[0110] Referring now to Fig. 20, which is an exemplary screen shot illustrating the creation of a security rule that restricts user CSingh from accessing all eRecords where the secure element 'item' (previously defined as XML element 'item\_no' in an Item Creation event per Fig. 13) has a value 4101. When the same query created in Fig. 19 is now executed on the evidence store by user CSingh, no eRecords are returned as shown in Fig. 21 because the security rule created in Fig. 20 prevents CSingh from viewing the eRecords. If user CSingh ran the query shown in Fig. 18, however, the results would be identical to those shown in Fig. 18 because the security rule created in Fig. 20 only applies to "item" elements and does not apply to "lot item" elements.

[0111] Having fully described several embodiments of the present invention, other equivalent or alternative methods of practicing the present invention will be apparent to those skilled in the art. For example, while system 10 was described as a distributed system, the system may be deployed in various other environments such as an enterprise environment, a stand-alone system, and the like. Also, while the present invention has been described using a particular combination of hardware and software, it should be recognized that other combinations of hardware and software are also



within the scope of the invention. Specifically, the invention may be implemented primarily in hardware, primarily in software, or using appropriate combinations thereof.

[0112] As another example, evidence store 30 was described as containing unstructured data in the form of XML documents. Embodiments of the invention can  
5 be used to access other types of unstructured data stored in the evidence store including, for example, data stored in other markup language formats. These and other embodiments as well as alternatives and equivalents to the invention will be recognizable to those of skill in the art after reading the description of the present invention. The scope of the invention should not, therefore, be determined solely by  
10 reference to the above description, but instead should be determined with reference to the appended claims along with their full scope of equivalents and alternatives.